Advanced Scientific Computing Research
FY 2006 Accomplishment

MHD Simulations on Massively Parallel (32,000 processors) Architectures

Paul Fischer and Fausto Cattaneo, Argonne National Laboratory

Summary

Magnetohydrodynamics (MHD) is central to many challenging physics problems of importance to the DOE Office of Science. We have developed a state-of-the-art code that features high-order numerical discretizations and multigrid solvers capable of scaling to thousands of processors to simulate MHD in complex domains.

MHD governs the motion and stability of many physical phenomena of interest to DOE, including astrophysical plasmas, geo- and solar dynamos, fusion plasmas, liquid metal cooling systems in nuclear reactors, and liquid metal plasma-facing material in tokamak side-walls and diverters. Most of these applications are dominated by convective transport and operate at high hydrodynamic and magnetic Reynolds numbers, Re and Rm, respectively. The governing physics is thus highly nonlinear and is essentially nondissipative. In addition, many of these applications involve complex domains that preclude the use of traditional global spectral methods.

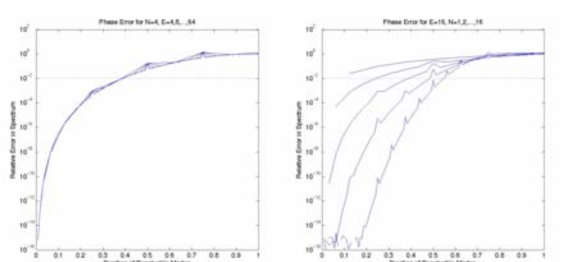

Figure 1. Phase error vs. fraction of resolvable modes for h-refinement (left) and p-refinement (right) for $u_t + u_x = 0$. The fraction of resolvable modes is increased only with increased order (p-refinement), which also yields rapid error reduction.

Argonne researchers have developed a numerical code for simulation of liquid metal MHD that couples state-of-the-art high-order numerical methods with the geometric flexibility required for the challenging MHD problems facing the DOE and scientific community.

Our MHD code is based on Argonne's hydrodynamics code Nek5000, which is a past Gordon Bell Prize winner that readily scales to thousands of processors. Nek5000 is based on the spectral element method (SEM), a high-order weighted residual technique that combines the geometric flexibility of finite elements with the rapid convegence and tensor-product efficiencies of global spectral methods. As with the finite element method, functions in the SEM are represented on compactly supported subdomains (elements), thereby simplifying implementation of complex boundary conditions. Grid refinement in the SEM is achieved by increasing the order of the polynomial representation within each element, with typical orders in the range of N= 8–16. The use of such high order minimizes numerical dissipation and dispersion and is important for high-Reynolds number applications where high wave-number error components are only

[*] Mathematics and Computer Science Division, (630) 252-6018, fischer@mcs.anl.gov

weakly damped by physical viscosity. This point is illustrated in Fig. 1, which shows that increasing approximation order leads to efficient use of computational resources. For modest error tolerances, a fivefold reduction in the number of gridpoints *per space dimension* is achievable by going from linear elements to 12^{th}-order elements.

In addition to minimizing MHD discretization errors, we have made significant strides in stabilized methods, on multigrid solvers, and porting to terascale platforms. Of particular importance to MHD is the development of dealiased quadrature rules that ensure energy conservation (Fig. 2). We have recently developed spectral element multigrid techniques that have proven to be two to three times faster than our earlier multilevel Schwarz methods across a range of applications.

A major advance this year has been the port of Nek5000 to the 32K processor Blue Gene platform at IBM Watson (BGW). BGW is a prototype of Petascale platforms anticipated to be in production in DOE within two years. The scalability of our approach is demonstrated by successful simulations of MHD problems involving a half-billion degrees-of-freedom requiring less than one second per semi-implicit timestep. Our project has been awarded three million node-hours, the first given as part of the BGW external science program, to study velocity distributions in turbulent MHD Taylor-Couette flows (Fig. 3) that serve as models for angular momentum transport in galaxies. Initial simulations, undertaken with support from a 2005 DOE INCITE award, indicated an unexpected linear angular velocity profile for this configuration. Higher-resolution studies on BGW are being used to confirm whether these trends persist at higher Reynolds numbers.

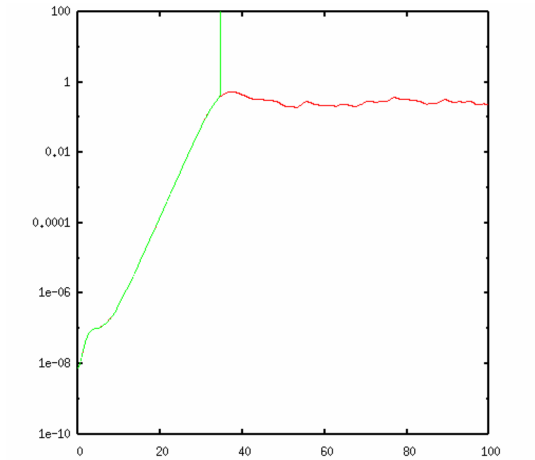

Figure 2: History of magnetic energy in an MHD benchmark for aliased (green) and dealiased (red) nonlinear evaluations. In the saturated nonlinear state, the aliased case is numerically unstable.

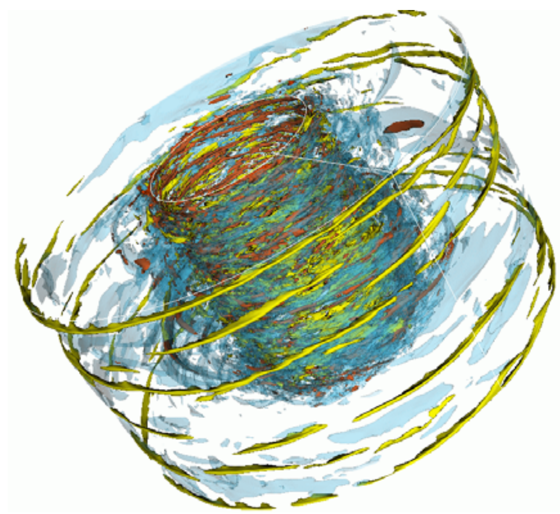
Figure 3: Turbulent MHD in Taylor-Couette flow at Re=62,000.

For further information on this subject contact:
Paul Fischer or Fausto Cattaneo
Argonne National Laboratory
fischer@mcs.anl.gov, cattaneo@mcs.anl.gov